

Micro-PSYPRE (Ψ_x^m): Toward Computing the Future

* Summary Report *

Selmer Bringsjord (PI)

Micah Clark

Trevor Housten

Konstantine Arkoudas (co-PI)

Department of Cognitive Science

Department of Computer Science

Rensselaer AI & Reasoning (RAIR) Laboratory:

<http://www.cogsci.rpi.edu/research/rair/>

Troy NY 12180 USA

selmer@rpi.edu

clarkm5@rpi.edu

houstt@rpi.edu

konstantine@alum.mit.edu

October 5, 2007

Report Documentation Page			Form Approved OMB No. 0704-0188		
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>					
1. REPORT DATE 05 OCT 2007	2. REPORT TYPE	3. DATES COVERED 00-00-2007 to 00-00-2007			
4. TITLE AND SUBTITLE Micro-PSYPRE(Ymx):Toward Computing the Future			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Department of Cognitive Science, Department of Computer Science,Rensselaer Polytechnic Institute,Carnegie 108, 110 8th Street,Troy,NY,12180			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 28	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Contents

1	Introduction & Background	3
1.1	The Problem	3
1.2	A Solution: The PSYPRE (Ψ) System	3
1.3	A Microcosm of the PSYPRE System: Micro-PSYPRE (Ψ_x^m)	5
1.4	Logic-based AI & Computational Cognitive Modeling	6
2	Electrifying Two Wargames	7
3	Micro-PSYPRE (Ψ_x^m) in Detail	12
3.1	Key Interrogatives	13
3.2	Representation Scheme	14
3.3	The “Frame Problem” Addressed	16
3.4	Planning	18
3.5	Post-diction & Projection	18
3.6	Demonstrations	20
4	The Future	22
5	Conclusion	25

1 Introduction & Background

The acute desire, on the part of the United States, to engage in Effects-Based Operations (or, for short, EBO; see e.g. Deptula 2001, Davis 2001) is here to stay. The name may of course change (and that's why we refer generically to *computing the future*), but the desire to determine what the future will bring, not just with respect to the physical effects of physical attack, but also with respect to the psychological effects of physical and psychological attack, will persist for a very long time.

Even our enemies understand that the exclusive use of classical, purely kinetic, force-on-force conflict is obsolete: It is certainly *not* the mode in which to take on the United States. Our enemies see their only chance in “unrestricted warfare” (Liang & Xiangsui 1999). If for no other reason than that we must fight fire with fire, and answer the various dimensions of asymmetrical conflict, we must rely on EBO, which by definition seeks to exploit the effects of blue force actions on individual and collective psyches of red and grey forces. Unfortunately, we must all confess on the other hand that, at least at present, EBO is just a pipe dream.

1.1 The Problem

What’s the source of the problem? That’s easy: vagueness. The kinetic effects of, say, a cruise missile can be calculated on the strength of clarity achieved by physics and engineering. But when trying to calculate the effects of such a missile on *morale*, clarity is hard to come by. EBO will never mature beyond fantasy unless we can formalize and mechanize the murky terms and concepts currently used to describe the cognitive side of it. In short, unless we have a formal account, with a corresponding implementation, of what it means to be a person, we will never be able to engage with any precision and reliability in EBO.

1.2 A Solution: The PSYPRE (Ψ) System

The solution to the problem is to engineer a system able to PREDict the effects, including the PSYchological effects, of actions and sequences of actions — we call this system PSYPRE or, for short, simply Ψ . Ψ is distinguished by eight attributes, which follow.

Attribute 1: Clear I/O Structure

Ψ will have a direct, straightforward *I/O* structure. It will take key queries from the human about the future, and quickly return concrete answers (accompanied by straightforward justifications).

Attribute 2: Transparent Virtual Environments

Ψ will be based on what we call *transparent* virtual environments: environments every inch of which are understood by the machine at a deep level. There will be no visual window dressing, information presented to the human user, but having no underlying formal representation that the machine can grasp and reason over.

Attribute 3: Next-Generation Virtual Humans

Ψ will take adversarial modeling to a revolutionary level because its simulation and prediction will be based upon virtual humans built as embodiments of the Rensselaer Advanced Synthetic Character Architecture for Logical Systems (RASCALS), which makes provision for robust epistemic and ethical reasoning not reachable by standard fare (e.g., SOAR). These virtual humans make it possible to predict the cognitive and behavioral effects of blue force actions on the adversary, and on surrounding population (grey forces).

Attribute 4: Based on Real Human Data

Ψ includes an implementation of a theory of human reasoning (Mental Meta-Logic; MML) based on data produced by state-of-the-art empirical investigation in the psychology of human reasoning.

Attribute 5: State-of-the-Art Machine Reasoning

Unprecedented predictive power will be provided by virtue of state-of-the-art machine reasoning systems that include cutting-edge automated reasoning, model building, and proof checking. These systems have been demonstrated to be effective in other projects (e.g., in the engineering of the Slate system, developed for ARDA/DTO's NIMD and A-SpaceX programs to enable hypothesis generation, testing, and analysis in the intelligence analysis domain).

Attribute 6: Predictions Accompanied by English Justifications

All Ψ 's predictions are accompanied by a justification, *and* that justification will be in plain English so that planners and analysts can readily understand the case being made by Ψ .

Attribute 7: Builds Upon Wargaming Technology

Ψ explicitly takes account of wargaming technology. Wargaming, beginning with H. G. Wells' (1913) *Little Wars*, has an illustrious history in the instruction of military strategists and tacticians, though few wargames have attempted to codify the *psychological* effects of players' actions. In work-

ing toward Ψ , we leverage wargames — particularly wargames of insurgent, revolutionary, and guerrilla warfare.

Attribute 8: Commercial-Grade Interface for Ease of Use

Ψ will offer human “pilots” of the system an interface that meets or exceeds the high standards of commercial games, through formative evaluation and usability testing.

To say PSYPRE is ambitious is an understatement, and we recognized “a full, frontal assault” was well beyond the scope of this small \$50k AFRL-sponsored effort.¹ So, we aimed instead at developing a microcosmic version of Ψ , micro-PSYPRE (Ψ_x^m), which can be instantiated for various strategy and wargames.

1.3 A Microcosm of the PSYPRE System: Micro-PSYPRE (Ψ_x^m)

An overview of the Ψ_x^m architecture is shown in Figure 1. The specific game of interest is indicated by replacing the subscripted variable x with the name of the game being referred to ($\Psi_{Nicaragua}^m$, for example, is an instantiation of the architecture for the specific game *Nicaragua!*, and is discussed later in §2). As this figure indicates, we can view Ψ_x^m as a function that takes in queries issued by planners and strategists interested in receiving predictions back. These predictions are all sought with respect to players in a wargame/strategy game that has a serious psychological dimension. Moreover, they are all sought within the context of at least a particular game state S_g , and possibly also within the context of the history (or partial history) of moves and states that led up to S_g . For a simple example, suppose that we want to know what player P_i is going to do if player P_j performs action a on his turn. Given this query as input, Ψ_x^m gives as output its prediction in the form of an action a' that it believes P_i will perform — *and* the output includes a clear justification for this prediction in the form of a deductive argument. We can also ask the system to make predictions given that *sequences* of actions are taken.

¹Our attack on the challenges involved in Ψ is multi-pronged, e.g., we have developed advanced linguistic capabilities (relevant to attributes 1 & 6) as part of our (ARDA/DTO-sponsored) Solomon project, a revolutionary QA system which possesses a *semantic* understanding of language.

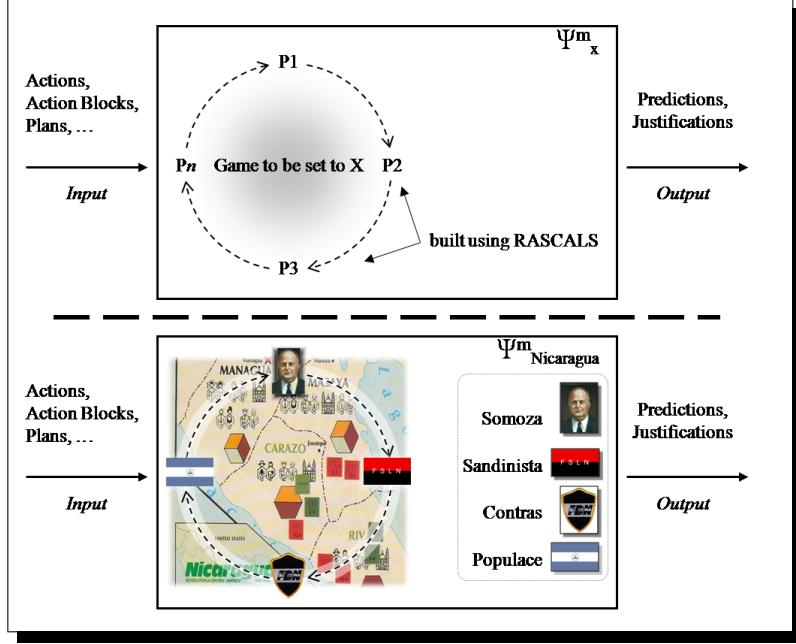


Figure 1: High-level architecture of Ψ_x^m and sample assignment of x to the psychologically-intense wargame *Nicaragua!*

1.4 Logic-based AI & Computational Cognitive Modeling

We intend for successive versions of Ψ_x^m to eventuate to Ψ . But, put bluntly, for the predictions made by Ψ to have any meaning, or to be of any value, attributes 3 & 4 must be accomplished; that is, the system must include, or at least facilitate, accurate cognitive models of our adversaries. In regard to Ψ_x^m , which is relegated to “games” of strategy and warfare, we can best enable a successful future Ψ by at present fastidiously employing the paradigm and formalisms of logic-based AI and logic-based computational cognitive modeling.²

The hallmarks of declarative (logic-based) computational cognitive modeling are the following two intertwined constraints: (1) The central units of information used in the approach are (at least in significant part) declarative in nature, and the central process carried out over these units is logical inference (deductive, inductive, abductive, analogical, etc.). (2) The approach to modeling the mind is top-down, rather than bottom-up. (These two points are interconnected because

²Otherwise, if winning “games” by any means were enough, we would simply use the techniques of traditional game-playing AI (e.g., game trees, min/max approximation, α - β search, etc.) and wait for Moore’s law to provide sufficient computational “horsepower” to play wargames perfectly.

once one commits to (1), (2) becomes quite unavoidable, since bottom-up processing in the brain, as reflected in today’s relevant formalisms (e.g., artificial neural networks), is based on units of information that are numerical, not declarative.) The systematization of declarative computational cognitive modeling is achieved by using formal logic, and hence declarative computational cognitive modeling, from the formal perspective, becomes logic-based computational cognitive modeling (LCCM).

How are logic-based computational cognitive modeling and human-level logic-based AI related? How similar are they? What makes them different? The encapsulated answer is straightforward: The two fields are largely based upon the same formalisms, both exploit the power of general-purpose programmable computing machines to process symbolic data, but LCCM targets computational simulations of human cognition, whereas human-level logic-based AI, as you might expect, strives to build beings that, at least behaviorally speaking, can pass for humans. While it is conceivable that both fields might well be on the same exact path (one that leads to building a computational system indistinguishable from a human), LCCM insists that the engineered system, at some suitably selected level of description, operate as a human does. Human-level AI would be content with artifacts that seem human, but under the hood really aren’t. As to shared formalisms, interested readers are directed to treatments of logic-based AI that introduce the relevant technical material (summarized e.g. in Bringsjord & Ferrucci 1998a, Bringsjord & Ferrucci 1998b, Nilsson 1991).

In this project, we directly followed the techniques and formalisms of LCCM, as set out and explained in (Bringsjord forthcoming).

2 Electrifying Two Wargames

To make Ψ_x^m concrete, we instantiated Ψ_x^m for the psychologically-intense wargame: *Nicaragua! Revolution in Central America* (Burtt & Miranda 1998). As indicated in Figure 2, *Nicaragua!* is a tabletop (i.e. “pencil and paper”) wargame. It portrays the conflicts between the Somoza, Sandinista and Contra factions in Nicaragua during the 1970s and ’80s. The game directly links military and political spheres, as it includes such decisive factors as intelligence operations and propaganda. In *Nicaragua!*, players assume the roles of the various warring government and rebel forces (the specific affiliations of government and rebel are dependent on the particular scenario being played out). Each side uses military forces to find and eliminate the enemy, as well as political forces to gain popular and foreign support, to recruit additional personnel, and demoralize the opposition.

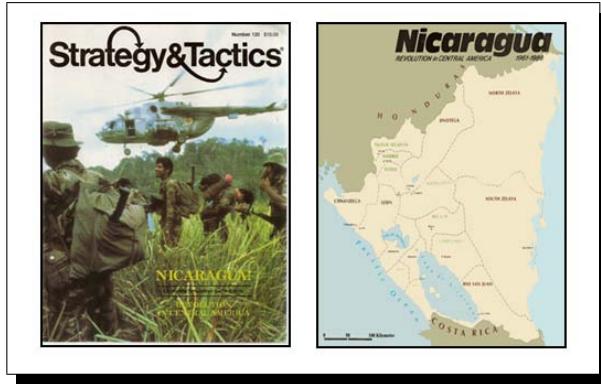


Figure 2: *Nicaragua! Revolution in Central America* tabletop wargame

Nicaragua! is played in a series of game-turns, each being a complete cycle through the following sequence.

1. World Events Phase: A chance occurrence of a major world event, e.g., a natural disaster, an independent assassination, a spontaneous revolt, a leader falls ill, or a major public relations/media disaster befalls a player.
2. Air Power Phase: Each player marshals and deploys their airborne forces.
3. Rebel Player's Turn: The following phases are performed (in order):
 - (a) Foreign Support Phase: Gains and/or losses of political and military support from neighboring countries and world powers are computed, and units of foreign forces are deployed to, and/or withdrawn from, the field.
 - (b) Political Phase:
 - i. Political Program Segment: A decision is made to keep the current political affiliation and treaties, or to adopt a new political program. The effects of changing political ideology, on national will and on the sympathies of the social classes and religious factions, are immediately computed and applied.
 - ii. Mobilization Segment: Political forces are used to recruit reinforcements, either as new units or as replacements for ongoing attrition. Alternatively, a mass mobilization may be declared, allowing units held in reserve to be deployed.

- iii. Political Warfare Segment: Political warfare is conducted. Also, leaders and emissaries may be dispatched to, and/or returned from, foreign countries.
 - (c) Organization Phase: Companies may be organized into battalions, or battalions broken down into companies. Units in separate locations may be organized into battle-groups, and the overt/covert status of units may be changed.
 - (d) Movement Phase: Units may be moved on the map, depending on the eligibility and properties of the units, terrain type, etc. The government player may declare reaction attacks against and searches of moving rebel forces.
 - (e) Reaction Phase (Government Player): Combat resulting from reaction attacks is resolved. The government player may deploy units held in reserve, and in addition, the government player may change the overt/covert status of his/her units.
 - (f) Intelligence Phase: Intelligence operations may be conducted in locations where there are covert government forces.
 - (g) Combat Phase: The rebel player may initiate combat in locations with qualifying rebel forces. The mode of combat, guerrilla or conventional, is selected by each player in secret before each battle is joined.
 - (h) Recovery Phase: Attempts may be made to remove neutralization effects from rebel forces. Also, one terrorized location may be checked for recovery.
 - (i) Nation Will Phase: The effects of previous phases on national will, social class support, and desertions, are computed and applied.
4. Government Player's Turn: The phases of the rebel player's turn are performed again, with the roles of rebel and government reversed.
5. Victory Determination: Each player checks to see if their scenario-dependent victory (and loss) conditions are met. Given the asymmetric nature of guerrilla warfare, it is possible for both players to win, and for both to lose. The specific victory and loss conditions may be kept private to each player — enhancing realism and increasing the difficulty of anticipating the opponent's strategies and actions.

Because of the nature of revolutionary warfare, there are a number of different tactics in *Nicaragua!* not usually found in conventional wargames;³ unfortunately,

³E.g., intentionally provoking harsh government suppression of an urban area riot, feigning a cease-fire to improve public relations in foreign countries and to rest troops, declaring martial law

we do not have room to cover them here. A sampling from the *Nicaragua!* rule-book is shown in Figure 3. Needless to say, *Nicaragua!* is not simple.



Figure 3: Two pages from the *Nicaragua!* rule-book

Our first order of business was to “electrify” *Nicaragua!*; that is, we implemented the entire tabletop version of *Nicaragua!* in a new, computerized game. A screen-shot from the electronic version of *Nicaragua!* is shown below in Figure 4. (The screen-shot is also hyper-linked to an on-line video demonstration of game-play, viz.: www.cogsci.rpi.edu/research/rair/wargaming/presentations/Nicaragua.wmv.) This electronic version supports Human vs. Human, Human vs. Computer, and Computer vs. Computer modes of play.

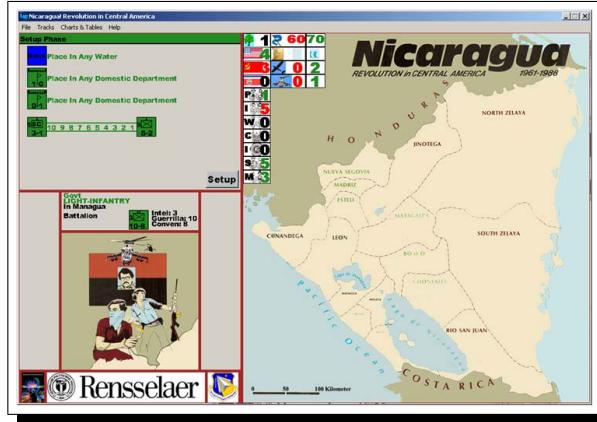


Figure 4: A screen-shot from the electronic version of *Nicaragua!*

and enforcing curfews in order to unmask covert rebel forces, repressing specific social classes, and using unrestricted firepower.

In addition to *Nicaragua!*, we developed our own wargame, *Iraq!*, which focuses on the political, religious, and economic factors of civil unrest in present-day, US-occupied Iraq. *Iraq!* is modeled directly on *Nicaragua!*, though in the adaptation, we have modified the rule-set to better suit the Iraqi scenario. The government forces in *Iraq!* represent a coalition of the US military, the Iraqi military, and the Iraqi police. The rebel forces are a mélange of militias, Al Qaeda, and Iran-supported foreign fighters. Within the game, the government needs to suppress sectarian violence and eliminate popular support for the various rebel forces, while avoiding direct confrontation with Iraq's meddling regional neighbors. The rebels' ultimate aims are to discredit the US-backed al-Maliki government, to foster all-out civil war, and to incite neighboring countries to invade and establish rule by *sharia* (Islamic law). A screen-shot of *Iraq!* is shown in Figure 5.

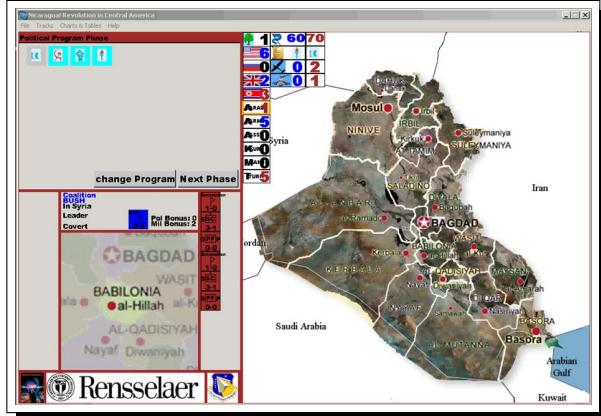


Figure 5: A screen-shot of *Iraq!*

Beyond being stand-alone computerized wargames, *Nicaragua!* and *Iraq!* function as game managers for Ψ_x^m (as shown in Figure 6, a second view of the Ψ_x^m architecture). The game manager performs three key functions within Ψ_x^m : (1) it mediates interactions between players (i.e. turns of the game) and enforces the game’s rules; (2) it publishes an audit trail of players’ observable actions (via APIs and log files); (3) it maintains the current game state (i.e. “the state of the world”) which can be inspected programmatically.

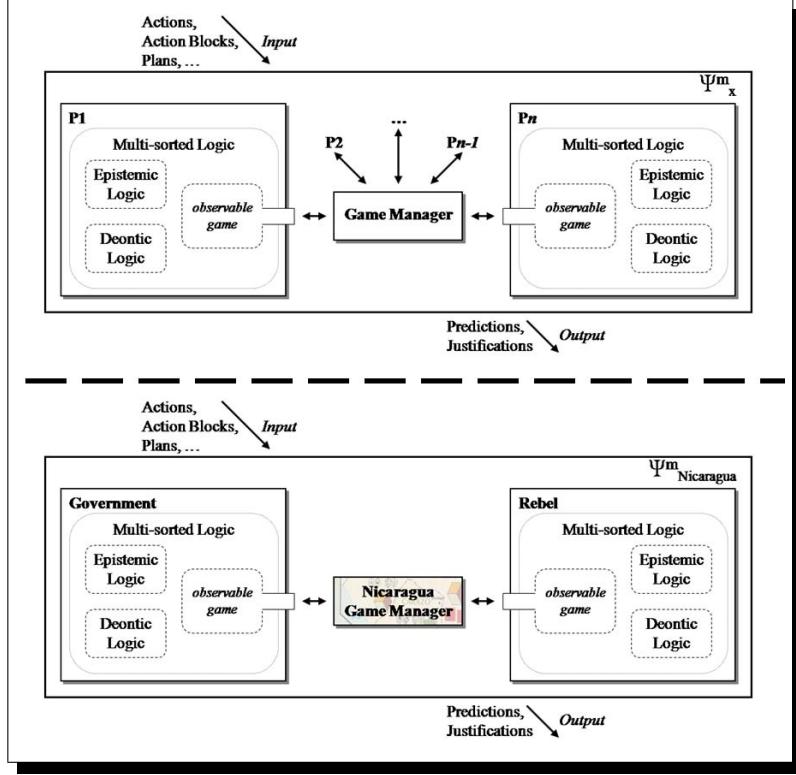


Figure 6: The game manager mediates interactions and enforces rules in Ψ_x^m

Nicaragua! and *Iraq!* are implemented entirely in Common Lisp, and they currently run on Mac OS-X and Microsoft Windows platforms. Initial synthetic characters are also implemented in Lisp; their reasoning capabilities (discussed later in §3.2–§3.5) are realized via the SNARK automated theorem prover (Stickel 2007), a multi-sorted, first-order logic theorem prover also implemented in Lisp. Both systems are functionally complete, but neither code-base is stable enough for external distribution.

3 Micro-PsyPRE (Ψ_x^m) in Detail

We have in mind three primary use-cases for Ψ_x^m , one for military and intelligence analysts, another for training military/intelligence students and professional wargamers, and a third for those playing only for entertainment.⁴ Of these three

⁴A fourth use-case is for adversary modelers, but we do not discuss this use-case here.

use-cases, the first two merit further discussion.

Military/Intelligence Analysts

Analysts seek predictions about particular scenarios, which they define. Defining a scenario involves specifying the current game state (i.e. “current state of the world”) and a (partial) history of players’ previous actions. Ψ_x^m is then run with only synthetic players participating (i.e. Computer vs. Computer), while actions and justifications are recorded (by Ψ_x^m) for later analysis. The analyst views the game-play from an external, third-person perspective. He/She may explore closely related variants of the main scenario (possibly varying the goals, strategies, and resources of the synthetic players as well as varying the game state and history) in a form of weighted Monte Carlo simulation.

Military/Intelligence Students & Professional Wargamers

Students and professional wargamers “test their metal” against computerized opponents (i.e. Human vs. Computer). Like analysts, students and professionals can define their own scenarios, if they wish. But students and professionals play the game from an embedded first-person perspective, i.e., they do not have access to the hidden, internal states of their opponents — they can only use “directly observable” information (provided by the game manager) from which to adduce opponent goals, strategies, and future actions. Students and professionals can avail themselves, in real-time, of some of the same powerful automated reasoning tools employed by synthetic players.

3.1 Key Interrogatives

Regardless of the use-case, Analyst or Wargamer, there are three key types of questions the user (or synthetic player) wishes to have answered about the current game: (1) “What has occurred and why?” (2) “What will occur and why?” (3) “How can my goals be achieved?” These three question types correspond to three types of reasoning queries: post-diction (or explanation), projection, and planning; a description of each type of reasoning follows.

Post-diction

Post-diction is the deducing of an explanation of the “current state of the world” from a *narrative*, a partial or incomplete history. In *Iraq!*, for example, if previously there were few insurgents in the As Suleymaniyah province, and now, for no obvious reason, there are a large number of insurgents, the user might ask: “What has occurred to increase insurgent forces in As Suleymaniyah?” Based on the narrative, Ψ_x^m reasons that either more foreign

fighters were smuggled in from northern Iran, or recruitment of local Sunni Kurds increased, or both. The Kurdish leadership's ardent claims to the oil fields (and oil revenue) of northern Iraq, and their non-Arab heritage, precludes increased support from Iran; while the Sunni Kurd's near-neutrality toward US occupation has, so far, kept low radicalization and recruitment to local militias. But, since insurgent strength has increased, something has changed. Based on observations (which are part of the narrative), two events are consistent with, and explain, the increase in insurgent forces: either (1) local imams have increased radicalization efforts among the poor and young, or (2) in response to concessions given by the US to the Shia, the local, middle-class Sunni have increased support for the militia, or both (1) & (2).

Projection

In projection, one supposes that a particular sequence of future actions will occur, then one asks “what will be the ramifications of this sequence of actions?” That is, the effects of the supposed actions are propagated forward as a prediction of the future world. For example, an analyst might ask “How will restoring water, power, and other basic services impact support for the Iraqi insurgency?” Ψ_x^m propagates forward the effects of restoring basic services, and determines that in rural areas and in areas with strong sectarianism, there will be little change in the insurgency’s strength; but in mixed urban areas — especially Baghdad, Mosul, and Basrah — support for Iraqi militias and foreign fighters decreases dramatically, and consequentially, violence and general instability are reduced.

Planning

Planning consists of generating a prescriptive, conditional sequence of actions which will accomplish or achieve some set of objectives. The objectives are specified as goals, constraints satisfied by the desired game state, and sub-goals. Intuitively, the analyst or wargamer is asking: “Given I desire X , what actions do I need to take in order to achieve X ?”. Plans can contain conditional sequences of actions, i.e., they contain contingency plans to countermand the uncertainty of opponent actions/reactions and of stochastic events.

3.2 Representation Scheme

The current representation scheme of Ψ_x^m is ostensibly the Situation Calculus (McCarthy & Hayes 1969, Levesque et al. 1998), a base set of predicates and functions for reasoning about progressions of game states (called *situations*) in multi-sorted, first-order logic. The situation calculus contains three sorts or types:

1. a sort for *situations* (i.e. game states) containing, at least, an initial situation distinguished by the constant S_0 ;
2. a sort for *actions*, things that cause or initiate change, denoted with variables a_0, a_1, \dots ;
3. a sort for *fluents*, things that can change over time or between situations, denoted with variables f_0, f_1, \dots , etc.

The situation calculus also contains the predicates ‘Holds’ and ‘Possible’, and the function ‘Result’:

$$\begin{aligned} \text{Holds} &: \text{fluent} \times \text{situation} \longrightarrow \{\text{true}, \text{false}\}; \\ \text{Possible} &: \text{action} \times \text{situation} \longrightarrow \{\text{true}, \text{false}\}; \\ \text{Result} &: \text{action} \times \text{situation} \longrightarrow \text{situation}. \end{aligned}$$

The ‘Holds’ predicate specifies that a fluent is true in a situation, e.g. $\text{Holds}(f_i, S_0)$ states that the fluent f_i is true in the initial situation; likewise, $\neg\text{Holds}(f_j, S_0)$ states that the f_j is false in the initial situation. Similarly, the ‘Possible’ predicate indicates that an action is, or is not, one of the permissible (legal) actions in a particular situation. Finally, the ‘Result’ function is used to denote the resultant situation arising from taking an action in a particular situation. Together, ‘Holds’ and ‘Result’ describe a fully-branching progression of situations, while ‘Possible’ thins the branches of impossible progressions — Figure 7 illustrates this branching structure.

These basic building-blocks of the situation calculus are used to specify a domain theory, a collection of axioms, specifically *action precondition* axioms, and *positive* and *negative effect* axioms. Action precondition axioms formalize the restriction that some set of conditions Π^i must be satisfied for action a_i to be permissible in a situation, e.g., it must be the Rebel player’s turn for any of the Rebel player’s actions to be possible. The action precondition axiom schema is:

$$\forall s [\Pi^i \rightarrow \text{Possible}(a_i, s)], \text{ where } s \text{ is in the sort of situations.}$$

Positive and negative effect axioms codify an action’s effect on a fluent — termed “positive” when it causes the fluent to hold in the resulting situation, and “negative” when it causes the fluent to no longer hold. Effects also have preconditions. For example, the positive effect of turning on a light-switch might be to cause the light to be on, a precondition for this would be that the light-bulb is not broken; if the bulb is broken, then the action of turning on the light-switch has no effect. We use $\Gamma^{i,j}$ and $\Phi^{i,j}$ to denote the set of positive and negative preconditions for action a_i to effect fluent f_j . The axiom schema for positive and negative effects are:

$$\begin{aligned} \forall s [(\text{Possible}(a_i, s) \wedge \Gamma^{i,j}) \rightarrow \text{Holds}(f_j, s)]; \\ \forall s [(\text{Possible}(a_i, s) \wedge \Phi^{i,j}) \rightarrow \neg\text{Holds}(f_j, s)]. \end{aligned}$$

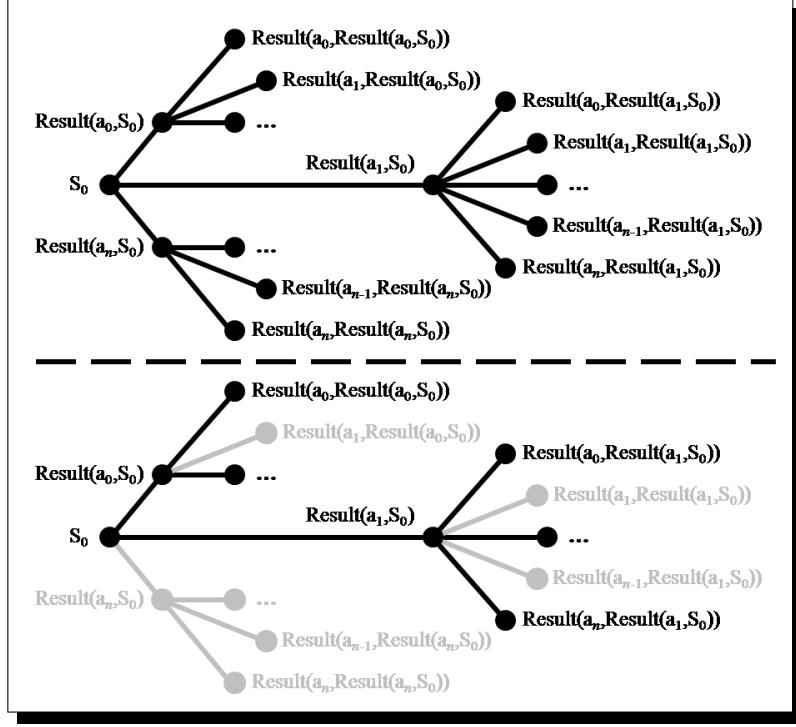


Figure 7: Branching structure of *situations* vs. *possible situations*

3.3 The “Frame Problem” Addressed

Effect axioms do not specify what fluents are *not* affected by actions. That is, the domain theory, as presented so far, does not allow us to infer that the values of unaffected fluents propagate unchanged into resulting situations. We need a way to stipulate e.g., that turning on a light-switch does not “turn off” the sun as well. Setting in place this common sense notion of inertia, that unaffected fluents propagate, is known as the “frame problem,” to which there are at least two generally accepted solutions: circumscription and frame axioms.

An explanation of circumscription would require more discourse and mathematical detail than is suitable here.⁵ We will with little discussion, however, give our reasons for not at present employing circumscription, a number of which are related to our need (due to budgetary and schedule constraints) to leverage, as much as possible, existing machine reasoning technologies, viz. automated the-

⁵ Interested readers are directed to Shanahan’s (1997b) *Solving the Frame Problem* for in-depth analysis of circumscription in the context of both the “frame problem” and the situation calculus.

orem provers such as SNARK, Otter, and Vampire. Our reasons, briefly stated, are:

- Circumscription requires second-order expressions. It cannot be properly represented in, or computed by, standard first-order theorem provers. So, use of circumscription would necessitate a meta-logical procedure to perform the minimization.
- The formula output produced by circumscription is not guaranteed to be expressible in first-order logic, i.e. even when all of the input formulas to circumscription are first-order sentences, the output may be second-order, and therefore, not usable by a standard first-order theorem prover⁶
- Circumscription (also called *predicate minimization*) is performed with respect to some set of predicates. As detailed by Shanahan (1997b), the answer returned for a variety of classic “frame problem” test cases (e.g., Yale shooting problem, Russian turkey shoot problem) depends in large part on what predicates are being minimized — with none giving acceptable answers on all problems. The frame axiom approach avoids this variability of success.
- Circumscription of first-order formulas, while expressible as a second-order formula, is really a semantic operation on the formal models of those first-order formulas. There is yet no clear theory as to how to incorporate semantic objects (e.g. models and counter-models) and semantic-based reasoning into formal arguments and linguistic (e.g. English) justifications.⁷

We address the “frame problem” by adopting the *successor-state* axiom schema from Reiter (1991). Successor-state axioms integrate positive and negative effect axioms and a commonsense law of inertia into a single axiom; the result being a single axiom per fluent. The successor-state axiom schema for a fluent f_j is:

$$\begin{aligned} \forall s \forall a [\text{Possible}(a, s) \rightarrow \\ (\text{Holds}(f_j, \text{Result}(a, s)) \leftrightarrow \\ \bigvee_{i=0}^n [\Gamma^{i,j} \wedge a = a_i] \\ \quad \vee \\ (\text{Holds}(f_j, s) \wedge \neg \bigvee_{i=0}^n [\Phi^{i,j} \wedge a = a_i]))], \end{aligned}$$

where $\Gamma^{i,j}$ and $\Phi^{i,j}$ are assumed false (unsatisfiable) for any action a_i not affecting fluent f_j . (A unique names axiom (UNA) is also needed, but we eschew that here.)

⁶There are ways to control the domain theory so as to guarantee that the result of circumscription will be first-order — but there are drawbacks there too (which are not discussed here).

⁷We are, at present, tacking the issue of integrating semantic objects into arguments as part of continued Slate development.

An example successor-state axiom for the hypothetical light-bulb would be:

$$\begin{aligned} \forall s \forall a [\text{Possible}(a, s) \rightarrow \\ (\text{Holds}(\text{On}(bulb), s), \text{Result}(a, s)) \leftrightarrow \\ ((\neg \text{Holds}(\text{Broken}(bulb), s) \wedge a = \text{TurnOn}(bulb)) \\ \vee \\ (\text{Holds}(\text{On}(bulb), s) \wedge a \neq \text{TurnOff}(bulb)))]. \end{aligned}$$

3.4 Planning

Planning is accomplished by using Green's (1969) method, a method of constructing goal-achieving plans as a side-effect of proving plan existence. To construct a plan, the following query is issued to a theorem prover:

$$\vdash ? \exists a (\text{Possible}(a) \wedge \text{Goal}(\text{Result}(a, S_0))),$$

where the ‘Goal’ predicate specifies the conditions to be achieved, i.e.,

$$\forall s (\text{Goal}(s) \leftrightarrow (condition1 \wedge condition2 \wedge \dots))$$

The representation scheme of Ψ_x^m is such that all (resolution-based) proofs of plan existence are constructive, i.e. a concrete plan of action can be directly extracted from the resulting proof of plan existence. The query above is for a single action (a single-step plan). The query is easily generalized for multi-step plans (action sequences) and conditional plans (see, e.g. Chapter 12 of Genesereth & Nissson 1987). Also, Green's (1969) method can be combined with traditional “goal decomposition” (i.e. sub-goaling) to construct and update strategic and tactical plans for synthetic players.

3.5 Post-diction & Projection

Post-diction and projection are both accomplished by a single common reasoning method wherein the actions occurring, and fluents holding, in a particular situation are constrained by observations of the actions and fluents in related situations. The central element of post-diction and projection is a *narrative*, a distinguished course of events about which the user may have incomplete information. The information that is known limits the possibilities for what is not known. In post-diction, the narrative is historical, and the “possibilities” deduced are what could/must have occurred. In projection, the narrative is of a supposed future, and the “possibilities” are what can/must occur.

We would like to specify narratives by stating what events have/will occur at specific points in time, say by using the predicate ‘Happens’, and by stating what

fluent held/will-hold at specific points in time using the predicate ‘HoldsAt’ (as exemplified in Figure 8). But, the situation calculus (as described in §3.2) has no facility for narratives, e.g., there is no way to distinguish the sequence of situations which has actually occurred from all other sequences that might have occurred but didn’t, and situations are *not* time-points.

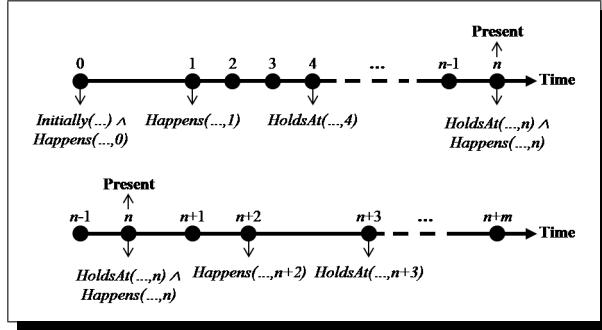


Figure 8: Narratives for Post-diction & Projection

Luckily, Miller & Shanahan (1994) showed how narratives could be gracefully added to the situation calculus by associating a situation with each time-point on a time-line. To do so, we introduce a new sort, *timepoint* (for convenience, we’ll assume that time-points correspond to natural numbers, and we assume the appropriate axioms for ordering ‘ $<$ ’ and addition ‘ $+$ ’) and a new function ‘State’:

$$\begin{aligned} \text{State} &: \text{timepoint} \longrightarrow \text{situation}, \text{ and} \\ \text{State}(0) &= S_0. \end{aligned}$$

With the function ‘State’ now in place, we can define a set of predicates suitable for specifying narratives:

$$\begin{aligned} \text{HoldsAt} &: \text{fluent} \times \text{timepoint} \longrightarrow \{\text{true}, \text{false}\}; \\ \text{Initially} &: \text{fluent} \longrightarrow \{\text{true}, \text{false}\}; \\ \text{Happens} &: \text{action} \times \text{timepoint} \longrightarrow \{\text{true}, \text{false}\}. \end{aligned}$$

The meaning of these new predicates are defined in terms of the initial situation calculus and the function ‘State’, as follows (where a , f , and t range over *actions*, *fluent*s, and *timepoints* respectively):

$$\forall f \forall t [\text{HoldsAt}(f, t) \leftrightarrow \text{Holds}(f, \text{State}(t))];$$

$$\forall f [\text{Initially}(f) \leftrightarrow \text{HoldsAt}(f, 0)];$$

$$\forall a \forall t [\text{Happens}(a, t) \leftrightarrow \text{State}(t + 1) = \text{Result}(a, \text{State}(t))].$$

The net result of integrating narratives (specifically the function ‘State’) into the situation calculus, is a reshaping (restricting) of the topology of situations into a tree structure (Figure 9), the trunk of which is the “known” narrative, and the branches are possible continuations of that narrative.⁸

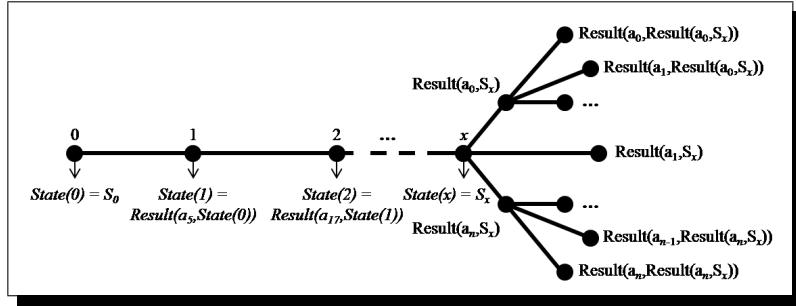


Figure 9: Branching structure of the Situation Calculus with Narratives

Once narratives are representable, post-diction and projection simply become queries about whether, given a narrative, certain things (collections of ‘Happens’ and ‘HoldsAt’ predicates) are possibly or necessarily true at particular time-points.

3.6 Demonstrations

We highlight Ψ_x^m (and the substantial progress toward Ψ which it represents) through three demonstration videos.⁹ The first video (Figure 10) showcases three use-cases for Ψ_x^m : the analyst, the professional wargamer, and the casual player.

The second demonstration video (Figure 11) shows an analyst who is interested in determining how the American-led coalition can have a stronger presence in Baghdad, Iraq. The analyst queries Ψ_x^m about the situation, and in return, is given a multi-step plan that accomplishes the analyst’s goal of a strengthened coalition presence.

The final demonstration (Figure 12) shows an analyst considering how best to disable Iran’s nuclear capability; via Ψ_x^m , several possible plans of action are developed. As the video suggests, we are extending *Iraq!* to cover (parallel) conflicts anywhere in the Middle East.¹⁰

⁸The specifics of the resultant topology are controlled by an existence-of-situations axiom, the

Current Modes of Play

Figure 10: A demonstration of three use-cases for Ψ_x^m
www.cogsci.rpi.edu/research/rair/wargaming/presentations/031007/MPPUseCases031007.wmv

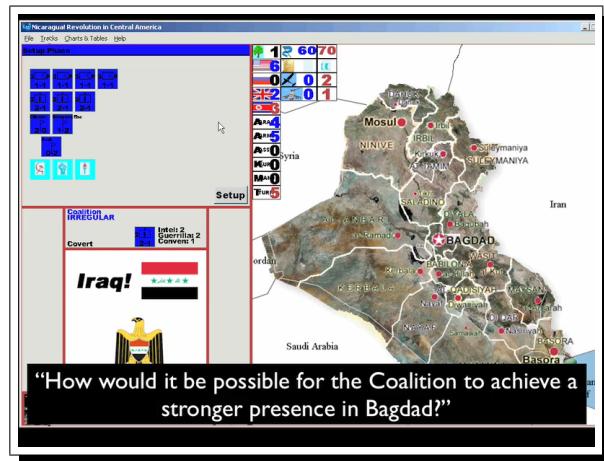


Figure 11: Planning, in Ψ_x^m , for a stronger coalition presence in Baghdad
www.cogsci.rpi.edu/research/rair/wargaming/presentations/031007/MPPMobilizeDemo031007.wmv



Figure 12: Analysis, in Ψ_x^m , for how to best disable Iran’s nuclear capability
www.cogsci.rpi.edu/research/rair/wargaming/presentations/031007/MPPVision031007.wmv

4 The Future

The project was, of course, quite small (\$50,000). It was designed to investigate whether a full development project aimed at Ψ makes sense to pursue. The result of investigation, by our lights, is that such pursuit makes good sense. But what, more specifically, would the next steps be? Here is a list of proposed next steps.

- **The representation scheme will be migrated to the event calculus.** Certain complex facets of commonsense reasoning (e.g., concurrent events, stochastic events, fluent trajectories, continuous change) are awkward to express in the situation calculus. The event calculus (perhaps best articulated in Mueller 2006) has a far richer predicate set than the situation calculus, and it directly addresses the problem of compactly stating the complex facets of commonsense reasoning.
- **The planning technique will be migrated to abductive planning.** We would have preferred to use the event calculus all along, but in keeping

ordering predicate and successor function for time-points, and (possibly) axioms of arboreality.

⁹Additional demonstration videos are available on the Ψ_x^m project’s website: www.cogsci.rpi.edu/research/rair/wargaming/.

¹⁰The successive wargame systems of *Nicaragua!*, *Iraq!*, and an extended *Iraq!* covering Iran, are leading us to a formal understanding of the general structures underlying games of strategic and tactical warfare. We are working toward a specification scheme for such games which is similar in intent to Stanford’s Game Description Language (GDL) (Love et al. 2006).

with our task-scoping principle that first-order theorem provers should be employed as much as possible, we could not do so. While the situation calculus allows planning to be performed via deduction (as explained in §3.4), the event calculus does not. Planning in the event calculus is accomplished by a technique known as “abductive planning” (described in detail e.g. in Shanahan 1997a, Shanahan 2000), wherein abductive logic programming is used to mirror the behavior of a partial-order hierarchical planner. Migration of the representation scheme to the event calculus (see the immediately preceding bulleted item) necessitates a migration to abductive planning as well.

- **We will develop a “cognitive event calculus.”** The particular calculi, situation or event, used to represent the physical world will be unimportant to EBO if one cannot equally express and propagate the cognitive “mental states” of adversaries and compatriots. We are working in earnest on a cognitive event calculus which will enable the representation and reasoning over the mental states of agents in multi-agent systems (we have already reported on some progress in this area, e.g. in Bringsjord et al. 2006). The “false-belief task” (Wimmer & Perner 1983) is one exemplar of reasoning about other agent’s beliefs. In the false-belief task, two people — person P_1 and person P_2 — are shown an object, which is then placed in one of two boxes. Person P_2 then leaves the area. While P_2 is gone (i.e. in the presence of P_1 alone) the object is moved from the first box to the other box. Person P_2 then returns, and P_1 is asked, “which box does P_2 believe the object to be in?” P_1 passes the test if he/she indicates the first box, and fails if he/she indicates the second box. To pass, P_1 must be able to understand that another’s mental representation of the world is different from their own. (The false-belief task has been used to detect the development of “perspective taking” in developmental psychology.) Below, in Figure 14, are two screen-shots (hyper-linked to corresponding movies) from our implementation of the false-belief task in the MMO virtual world *Second Life*. The first screen-shot (and corresponding movie) is of an immature subject who fails the false-belief task; the second screen-shot (and movie) are to a mature subject who is able to correctly reason about other’s beliefs.
- **The “frame problem” solution will be migrated to circumscription.** While in §3.3 we gave objections to the use of circumscription, there are several reasons for using circumscription over successor-state axioms. First, circumscription simplifies the domain theory — not in total number of axioms needed (which is largely unaffected), but in the complexity of the in-



Figure 13: Failing the “false-belief task” in Second Life
www.cogsci.rpi.edu/research/rair/asc_rca/20070810.SL.Cog.FalseBelief.Failure.wmv



Figure 14: Passing the “false-belief task” in Second Life
www.cogsci.rpi.edu/research/rair/asc_rca/20070629.SL.Noncog.FalseBelief.Pass.wmv

dividual axioms. A more weighty reason for using circumscription is that, when properly constrained, the minimal models produced by circumscription are equivalent to the “partial mental models” of Mental Model Theory (Johnson-Laird 2006) and Mental Meta-logic (Yang & Bringsjord 2001, Yang & Bringsjord 2003), two psychological theories that accurately predict and explain systematic errors in human reasoning (so called “cognitive

illusions”) which have been empirically observed in a number of domains (e.g., decision-making, game-playing, and deductive, probabilistic, & deontic reasoning, some of which are surveyed in Johnson-Laird & Yang 2006).¹¹ Through circumscription, we can simultaneously realize prescriptive reasoning (what a human agent/opponent ought to believe and do) and predictive reasoning (what a human agent/opponent will likely believe and do) in Ψ_x^m .

5 Conclusion

As mentioned, this was a small proof-of-concept project. We do conclude that the concept in question has, in fact, been validated: it is possible to augment the decision-making and future-determining power of humans by computing the psychological future in asymmetrical conflict. Realizing this possibility is our objective, given sufficient investment from sponsors.

References

- Bringsjord, S. (forthcoming), Declarative/Logic-Based Computational Cognitive Modeling, in R. Sun, ed., ‘The Cambridge Handbook of Computational Psychology’, Cambridge University Press, Cambridge, UK. Preprint available at http://kryten.mm.rpi.edu/sb_lccm_ab-toc_031607.pdf.
- Bringsjord, S., Arkoudas, K. & Yang, Y. (2006), Final Report on: “New Architectures, Algorithms, and Designs that Lead to Implemented Machine Reasoning over Knowledge in Epistemic & Deontic Formats, in the Service of Advanced Wargaming”, technical report, Rensselaer AI & Reasoning Laboratory (RAIR), Rensselaer Polytechnic Institute (RPI), Troy, NY. A report to AFRL about Agreement No. FA8750-04-01-1065.
- Bringsjord, S. & Ferrucci, D. (1998a), ‘Logic and artificial intelligence: Divorced, still married, separated...?’, *Minds and Machines* **8**, 273–308.
- Bringsjord, S. & Ferrucci, D. (1998b), ‘Reply to Thayse and Glymour on logic and artificial intelligence’, *Minds and Machines* **8**, 313–315.
- Burtt, J. & Miranda, J. (1998), ‘Nicaragua! Revolution in Central America’, *Strategy & Tactics* **120**.

¹¹It is also worth noting that “cognitive illusions,” of the type explained by partial mental models, are independent of language and culture (Yang et al. 2005).

- Davis, P. K. (2001), *Effects-Based Operations (EBO): A Grand Challenge for the Analytical Community*, RAND. ISBN: 0-8330-3108-2; MR-1477-USJFCOM/AF.
- Deptula, D. S. (2001), *Effects-Based Operations: Changes in the Nature of Warfare*, Aerospace Education Foundation, Arlington, VA.
- Genesereth, M. R. & Nilsson, N. J. (1987), *Logical Foundations of Artificial Intelligence*, Morgan Kaufmann, Palo Alto, CA.
- Green, C. (1969), Application of theorem proving to problem solving, in ‘Proceedings of the First International Joint Conference on Artificial Intelligence (IJCAI-69)’, IJCAII, Washington, DC, pp. 219–239.
- Johnson-Laird, P. N. (2006), *How We Reason*, Oxford University Press, Oxford, UK.
- Johnson-Laird, P. N. & Yang, Y. (2006), Mental logic, mental models, and computer simulations of human reasoning, in R. Sun, ed., ‘Cambridge Handbook of Computational Cognitive Modeling’, Cambridge University Press, Cambridge, UK. In press.
- Levesque, H., Pirri, F. & Reiter, R. (1998), ‘Foundations for the Situation Calculus’, *Linköping Electronic Articles in Computer and Information Science* 3(nr 18). <http://www.ep.liu.se/ea/cis/1998/018/>.
- Liang, Q. & Xiangsui, W. (1999), *Unrestricted Warfare*, PLA Literature and Arts Publishing House, Beijing, China.
- Love, N., Hinrichs, T. & Genesereth, M. (2006), General Game Playing: Game Description Language Specification, technical report LG-2006-01, Stanford Logic Group, Computer Science Department, Stanford University, Stanford, CA.
- McCarthy, J. & Hayes, P. J. (1969), Some philosophical problems from the standpoint of artificial intelligence, in B. Meltzer & D. Michie, eds, ‘Machine Intelligence 4’, Edinburgh University Press, Edinburgh, Scotland, pp. 463–502.
- Miller, R. & Shanahan, M. (1994), ‘Narratives in the Situation Calculus’, *The Journal of Logic and Computation* 4(5), 513–530.
- Mueller, E. T. (2006), *Commonsense Reasoning*, Morgan Kaufmann, San Francisco, CA.

- Nilsson, N. (1991), ‘Logic and Artificial Intelligence’, *Artificial Intelligence* **47**, 31–56.
- Reiter, R. (1991), The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression, in V. Lifschitz, ed., ‘Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy’, Academic Press, San Diego, CA, pp. 359–380.
- Shanahan, M. (1997a), Event Calculus Planning Revisited, in S. Steel & R. Alami, eds, ‘Recent Advances in AI Planning, Proceedings of the 4th European Conference on Planning (ECP’97)’, Vol. 1348 of *Lecture Notes in Computer Science*, Springer, Berlin, Germany, pp. 390–402.
- Shanahan, M. (1997b), *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*, MIT Press, Cambridge, MA.
- Shanahan, M. (2000), ‘An abductive event calculus planner’, *The Journal of Logic Programming* **44**(1–3), 207–240.
- Stickel, M. E. (2007), ‘SNARK – SRI’s New Automated Reasoning Kit’. <http://www.ai.sri.com/~stickel/snark.html>.
- Wells, H., Gygax, G., Varhola, M. & Varhola, D. (2004/1913), *Little Wars*, Skrimisher Publishing, Alexandria, VA.
- Wimmer, H. & Perner, J. (1983), ‘Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children’s understanding of deception’, *Cognition* **13**(1), 103–128.
- Yang, Y. & Bringsjord, S. (2001), Mental Metalogic: A New Paradigm for Psychology of Reasoning, in L. Chen & Y. Zhou, eds, ‘Proceedings of the Third International Conference of Cognitive Science (ICCS 2001)’, USTC Press, Hefei, China, pp. 199–204.
- Yang, Y. & Bringsjord, S. (2003), Mental metalogic and its empirical justifications: The case of reasoning with quantifiers and predicates, in R. Alterman & D. Kirsh, eds, ‘Proceedings of the 25th Annual Conference of the Cognitive Science Society (CogSci 2003)’, Cognitive Science Society, Boston, MA, pp. 1275–1280.
- Yang, Y., Zhao, Y., Zeng, J., Guo, J., Ju, Z. & Bringsjord, S. (2005), Empirical Justifications for the Universalness of the Mental Logic and Mental Models Paradigm, in B. G. Bara, L. Barsalou & M. Bucciarelli, eds, ‘Proceedings of

the 27th Annual Conference of the Cognitive Science Society (CogSci 2005)', Lawrence Erlbaum, Mahwah, NJ, pp. 2399–2403.